

Distributionally Robust Safe Control of Robotic Manipulators in Dynamic Environments

Abstract—In this paper, we investigate safe execution for robotic manipulators operating in environments with dynamic obstacles and perception uncertainty. A key challenge in this setting is that obstacle states must be inferred from noisy measurements, and estimation errors can render safety constraints overly optimistic, increasing the risk of violations. To address this, we propose a distributionally robust control framework that integrates Kalman-filter-based obstacle estimation with control barrier function (CBF) safety constraints, explicitly accounting for estimation uncertainty in safety-critical control. Using a Kalman filter, we maintain a Gaussian belief over obstacle positions and velocities, construct an ambiguity set capturing plausible deviations, and derive a robust CBF constraint that enforces safety under the worst case within this set. We validate the effectiveness and robustness of our approach through simulation studies, demonstrating safe manipulation under uncertainty and dynamic obstacles, and compare its performance against baseline methods such as standard point-estimate CBF controllers and fixed-margin safety constraints. The codes and videos are available at our project page: <https://robotic-manipulators.github.io/DR-CBF/>.

I. INTRODUCTION

Ensuring safety with formal guarantees remains a fundamental challenge for modern robotic manipulators. This problem becomes particularly difficult when robots operate in shared and dynamic environments, where they must provide strict real-time safety guarantees while executing high-frequency control in the presence of unpredictable moving obstacles. Recently, control barrier functions (CBFs) [2] have emerged as a promising approach for formally certifying safety in such scenarios. As proactive safety filters, CBFs can be integrated into a quadratic programming framework that minimally modifies a nominal control input to enforce strict safety constraints online. By mathematically certify the forward invariance of desired safe sets, CBFs provide both rigorous safety guarantees and real-time applicability [1].

Most existing CBF-based methods typically assume exact knowledge of the robot’s state and precise information about the surrounding environment. However, in open and interactive settings, such information is usually not available a priori, and manipulators must instead infer the states of dynamic obstacles from noisy sensor measurements. Compounded errors from perception and state estimation can make nominal safety constraints overly optimistic, potentially leading to safety violations [8], [6]. While recent studies have explored robust and probabilistic CBF formulations to explicitly address these uncertainties [7], [23], many practical implementations still rely on point estimates from filters or heuristic, fixed-margin safety buffers. Such approaches can be unsafe under high noise or overly conservative, limiting

the manipulator’s operational efficiency.

Motivated by the limitations of existing CBF-based approaches under noisy observations, we propose a novel distributionally robust control framework that explicitly accounts for estimation uncertainty in safety-critical control. Our approach directly integrates Kalman-filter-based obstacle estimation with CBF safety constraints. Using a Kalman filter, we maintain a Gaussian belief over obstacle positions and velocities based on noisy sensor measurements. We then construct an ambiguity set that captures plausible deviations of the dynamic obstacle states and derive a robust CBF constraint that enforces safety under the worst-case scenario within this set, leveraging principles from distributionally robust optimization [9].

The main contributions of this paper are as follows:

- First, we develop a distributionally robust control framework for robotic manipulators that explicitly integrates Kalman-filter-based obstacle estimation with safety constraints to account for perception uncertainty. This framework ensures that safety-critical control actions remain valid even under noisy measurements and estimation errors, providing formal safety guarantees in dynamic and interactive environments.
- Second, from a technical perspective, we construct an ambiguity set derived from the Gaussian belief over obstacle positions and velocities obtained from the Kalman filter. This allows us to systematically capture plausible deviations of dynamic obstacles and derive worst-case robust CBF constraints, rigorously enforcing safety for all distributions within the ambiguity set and extending the standard CBF framework to handle probabilistic perception uncertainty.
- Finally, we implement the proposed approach and validate its effectiveness and robustness in simulation studies. The results demonstrate superior safe manipulation performance in dynamic environments compared to standard point-estimate CBF controllers and fixed-margin safety baselines, highlighting the practical benefits of explicitly accounting for estimation uncertainty in real-time robotic control.

II. RELATED WORKS

Motion planning for robotic manipulators with collision avoidance has a long research history. In dynamic environments, methods have been developed to predict moving obstacles and filter unsafe actions [10], [20]. Accurate environment representation is crucial for evaluating safety constraints; for example, signed distance functions (SDFs)

provide distance and gradient information essential for reactive control [16]. More recently, neural implicit representations have been used for complex manipulator geometries and joint-space collision avoidance [11], [13]. However, constructing these representations at high control frequencies remains challenging under perception uncertainty and noisy sensor measurements.

Distributionally robust optimization (DRO) provides a principled framework for handling parameter uncertainty, particularly when the true probability distribution is unknown or only partially observable [9]. DRO compensates for discrepancies between empirical and true distributions by optimizing over an ambiguity set, such as those defined by moments [22] or Wasserstein distances [24]. Recently, DRO has been applied in robotics and control to provide robust performance guarantees against distributional shifts [5], [17], [19], with applications including Wasserstein-safe RRT for path planning [12], robust MPC using Wasserstein tubes [3], and safe control synthesis under model uncertainty [14].

Our work is motivated by recent applications of DRO for safe mobile robot navigation, which explicitly account for sensor and state estimation uncertainties [15]. However, while their approach relies on finite sensory samples to construct an empirical distribution and solves the resulting DRO problem in the dual space, robotic manipulators considered in our work introduce additional challenges. In particular, obtaining sufficient samples for rapidly moving dynamic obstacles is often impractical. To address this, we use a Kalman filter to maintain a Gaussian belief over each obstacle's position and velocity, constructing the ambiguity set directly from this distribution. By leveraging the structural properties of Gaussian distributions, our method avoids computationally intensive dual-space reformulations. We approximate the ambiguity set and solve the robust safety constraints directly, ensuring computational efficiency while supporting the high-frequency control rates required for safe dynamic manipulation.

III. PROBLEM FORMULATION

We consider a manipulator whose configuration evolves according to a control-affine model

$$\dot{x} = f(x) + g(x)u \quad (1)$$

where $x(t) \in \mathbb{R}^n$ denotes the system state; $u(t) \in U \subset \mathbb{R}^m$ is the control input; $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $g: \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ are two continuous functions.

The robot operating in the dynamic environments with moving obstacles. Let $\xi(t) \in \mathbb{R}^n$ denote the real-time position of the dynamic obstacle and $v_\xi(t) = \dot{\xi}(t)$ be its velocity. Safety requires that $x(t)$ maintain a minimum distance $d_s > 0$ from the obstacle at all times. Formally, we define the safety function as

$$h(x(t), \xi(t)) = \|x(t) - \xi(t)\|_2^2 - d_s^2, \quad (2)$$

Then the time-varying safe set is defined as

$$\mathcal{C}(t) = \{x \in \mathbb{R}^n : h(x, \xi) \geq 0\}. \quad (3)$$

In this paper, we adopt control barrier function (CBF) [2] to enforce safety. Specifically, the system remains within the safe set if there exists a class \mathcal{K} function α such that the following CBF constraint is satisfied for all $t \geq 0$,

$$\frac{\partial h}{\partial x}(f(x) + g(x)u) + \frac{\partial h}{\partial \xi}v_\xi(t) \geq -\alpha(h(x(t), \xi(t))) \quad (4)$$

To facilitate the integration of this constraint into an optimization framework, we denote the above constraint in the standard form $G(\cdot) \leq 0$ as:

$$\begin{aligned} & -\frac{\partial h}{\partial x}(f(x) + g(x)u) - \frac{\partial h}{\partial \xi}v_\xi(t) - \alpha(h(x(t), \xi(t))) \\ & =: G(u(t), \xi(t), v_\xi(t)) \leq 0 \end{aligned} \quad (5)$$

The exact obstacle state is often unavailable in practice due to sensor noise and environmental disturbances. In particular, obstacle velocity is difficult to measure directly. To account for this uncertainty, we model obstacle motion as a stochastic process, summarized by the following assumption.

Assumption 1. *The distribution of the obstacle's state $(\xi(t), v_\xi(t))$ follows an unknown Gaussian distribution, i.e.,*

$$(\xi(t), v_\xi(t)) \sim \mathcal{N}(\mu(t), \Sigma(t)) =: \mathbb{P}_\xi(t) \quad (6)$$

where $\mu(t)$ is the mean vector of the state and $\Sigma(t)$ is the associated covariance matrix.

To this end, the CBF constraint $G(u(t), \xi(t), v_\xi(t)) \leq 0$ is then replaced by the chance constraint, which ensures that the system remains safe with a high probability. In particular, for each time instant t , we require that

$$\mathbb{P}_\xi(t) (G(u(t), \xi(t), v_\xi(t)) \leq 0) \geq 1 - \epsilon \quad (7)$$

where $\epsilon \in (0, 1)$ is the risk parameter.

However, in real-world applications, the true distribution $\mathbb{P}_\xi(t)$ of the obstacle state is generally unknown. To capture this distributional uncertainty, we denote by $(\hat{\xi}(t), \hat{v}_\xi(t))$ an estimate of the obstacle state obtained by some state estimation technique, and construct an ambiguity set $\mathbb{D}(t)$ based on this estimate that contains the true obstacle-state distribution, i.e., $\mathbb{P}_\xi(t) \in \mathbb{D}(t)$. Based on this, we formulate the problem studied in this paper as follows.

Problem 1 (Distributionally Robust Safe Control). *Given a manipulator modeled by system (1) operating in a dynamic environment where the obstacle is described by $(\xi(t), v_\xi(t))$, and a safety function $h(x(t), \xi(t))$ defined in (2), under Assumption 1, synthesize a control input $u(t)$ such that*

$$\inf_{\mathbb{P}(t) \in \mathbb{D}(t)} \mathbb{P}(G(u(t), \xi(t), v_\xi(t)) \leq 0) \geq 1 - \epsilon.$$

IV. MAIN RESULTS

In this section, we first employ a Kalman filter to infer the obstacle's state from noisy measurements, yielding a Gaussian belief, and then construct an ambiguity set that captures plausible deviations of the obstacle's distribution. Finally, we develop a tractable convex reformulation by introducing a surrogate ambiguity set that enables a computationally efficient approximation amenable to real-time control.

A. Kalman Filtering for Dynamic Obstacle Estimation

To infer the obstacle motion from noisy perception, we employ a Kalman filter for the state estimation. This is motivated by the fact that obstacle's position can often be measured reliably even in the presence of noise, whereas its velocity is typically not directly observable and must be inferred from temporal position changes. The Kalman filter provides a lightweight estimator that recursively fuses a simple motion model with position measurements, yielding both a velocity estimate and a quantified uncertainty.

The update of the Kalman filtering is given as follows. At each time instant t , given the prior belief $(\hat{\mu}^-(t), \hat{\Sigma}^-(t))$ and a new measurement $y(t) = \xi(t) + n(t)$ where $n(t)$ is the measurement noise, the filter performs a measurement update by computing the Kalman gain

$$K(t) = \hat{\Sigma}^-(t)H^\top \left(H\hat{\Sigma}^-(t)H^\top + R \right)^{-1},$$

and updating the posterior mean and covariance as

$$\begin{aligned} \hat{\mu}(t) &= \hat{\mu}^-(t) + K(t)(y(t) - H\hat{\mu}^-(t)), \\ \hat{\Sigma}(t) &= (I - K(t)H)\hat{\Sigma}^-(t), \end{aligned}$$

where $H = I$ denotes the measurement matrix and R is the measurement noise covariance. The prior $(\hat{\mu}^-(t), \hat{\Sigma}^-(t))$ is obtained from the motion model through the standard prediction step, with process noise covariance Q .

As a result, conditioned on the measurement history up to time t , the estimated obstacle state admits a Gaussian belief

$$(\hat{\xi}(t), \hat{v}_\xi(t)) \sim \mathcal{N}(\hat{\mu}(t), \hat{\Sigma}(t)) =: \hat{\mathbb{P}}(t),$$

where $\hat{\mu}(t)$ and $\hat{\Sigma}(t)$ are the estimated mean and covariance, respectively.

B. Distributionally Robust Safety Constraint

While the Kalman Filter provides an estimate $\hat{\mathbb{P}}(t)$, there often exists a ‘‘reality gap’’ due to estimation errors. Therefore, relying solely on $\hat{\mathbb{P}}(t)$ may lead to overconfident control actions and potential safety violations when the true distribution deviates from the estimate. To account for this distributional mismatch, we construct an ambiguity set that collects all plausible distributions around $\hat{\mathbb{P}}(t)$. In this paper, we quantify the discrepancy between distributions via the 2-Wasserstein distance, as it provides a geometrically meaningful metric that compares distributions in the state space. The 2-Wasserstein distance is defined as follows.

Definition 1 (2-Wasserstein Distance). *Given two distributions $\mathbb{P}_1, \mathbb{P}_2$, 2-Wasserstein distance between \mathbb{P}_1 and \mathbb{P}_2 is*

$$\mathbb{W}(\mathbb{P}_1, \mathbb{P}_2) = \left(\inf_{\pi \in \Pi(\mathbb{P}_1, \mathbb{P}_2)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|\zeta_1 - \zeta_2\|_2^2 \pi(d\zeta_1, d\zeta_2) \right)^{\frac{1}{2}} \quad (8)$$

where $\Pi(\mathbb{P}_1, \mathbb{P}_2)$ denotes the set of all couplings, that is, all joint distributions of the random variables ζ_1 and ζ_2 with marginal distributions \mathbb{P}_1 and \mathbb{P}_2 , respectively.

Based on this metric, we define our ambiguity set around the belief distribution $\hat{\mathbb{P}}(t)$ as a Wasserstein ball.

Definition 2 (Ambiguity Set). *Given the estimated distribution of the dynamic obstacle $(\hat{\xi}(t), \hat{v}_\xi(t)) \sim \hat{\mathbb{P}}(t)$, we define the ambiguity set around $\hat{\mathbb{P}}$ as*

$$\mathbb{D}(t) = \{ \mathcal{N}(\mu(t), \Sigma(t)) : \mathbb{W}(\hat{\mathbb{P}}(t), \mathcal{N}(\mu(t), \Sigma(t))) \leq r(t) \}, \quad (9)$$

where r is a radius to be chosen for the Wasserstein ball.

In this work, we dynamically scale the radius based on the filter's own uncertainty. Specifically, we choose

$$r(t) = \kappa \cdot \sqrt{\text{tr}(\hat{\Sigma}(t))}, \quad (10)$$

where κ is a tuning parameter.

Then, by considering the worst-case distribution within $\mathbb{D}(t)$, we arrive at the following **distributionally robust optimization** (DRO) problem:

$$u^*(t) = \arg \min_{u(t) \in \mathbb{R}^m} \|u(t) - u_{\text{norm}}(t)\|_2^2 \quad (11a)$$

$$\text{s.t.} \quad \inf_{\mathbb{P} \in \mathbb{D}(t)} \mathbb{P}(G(u(t), \xi(t), v_\xi(t)) \leq 0) \geq 1 - \epsilon, \quad (11b)$$

where $u_{\text{norm}}(t)$ denotes a nominal controller that achieves the desired task objective in the absence of safety constraints. The objective seeks minimally invasive safe control input by staying as close as possible to $u_{\text{norm}}(t)$, while the constraint enforces a probabilistic safety guarantee against the worst-case distribution in $\mathbb{D}(t)$ with risk level ϵ .

As a result, if the radius κ is chosen such that the true distribution $\mathbb{P}_\xi(t)$ is contained within the ambiguity set $\mathbb{D}(t)$, the resulting controller $u^*(t)$ ensures the safety guarantee (7).

C. Tractable Convex Reformulation

The DRO constraint in (11) is in general computationally intractable, since the 2-Wasserstein ball around a Gaussian belief couples the mean and covariance through a non-separable Bures-type term. To enable real-time control, we derive a tractable reformulation by constructing a *surrogate ambiguity set* that over-approximates the original Wasserstein ambiguity set. The key idea is to decouple the uncertainty in the mean and the covariance by bounding the mean and the covariance eigenvalues within a spectral interval, which yields a convex set amenable to analytical bounding.

Formally, given belief distribution $\hat{\mathbb{P}}(t) = \mathcal{N}(\hat{\mu}(t), \hat{\Sigma}(t))$, we define the surrogate ambiguity set $\mathbb{D}_{\text{sur}}(t)$ as

$$\mathbb{D}_{\text{sur}}(t) = \{ \mathcal{N}(\mu, \Sigma) : \|\mu - \hat{\mu}\|_2 \leq r, \underline{\lambda}I \preceq \Sigma \preceq \bar{\lambda}I \} \quad (12)$$

where the spectral bounds are given by

$$\underline{\lambda}(t) = \left(\max \left\{ 0, \sqrt{\lambda_{\min}(\hat{\Sigma}(t))} - r(t) \right\} \right)^2, \quad (13)$$

$$\bar{\lambda}(t) = \left(\sqrt{\lambda_{\max}(\hat{\Sigma}(t))} + r(t) \right)^2. \quad (14)$$

The following proposition guarantees that $\mathbb{D}_{\text{sur}}(t)$ is a conservative relaxation of $\mathbb{D}(t)$, and thus replacing $\mathbb{D}(t)$ with $\mathbb{D}_{\text{sur}}(t)$ preserves the desired worst-case safety guarantee.

Proposition 1. Let $\mathbb{D}(t)$ be the 2-Wasserstein ball centered at the Gaussian belief $\hat{\mathbb{P}}(t) = \mathcal{N}(\hat{\mu}(t), \hat{\Sigma}(t))$, i.e.,

$$\mathbb{D}(t) = \left\{ \mathcal{N}(\mu, \Sigma) : \mathbb{W}_2(\mathcal{N}(\mu, \Sigma), \hat{\mathbb{P}}(t)) \leq r(t) \right\}.$$

Then the surrogate ambiguity set satisfies

$$\mathbb{D}(t) \subseteq \mathbb{D}_{sur}(t).$$

Proof. Consider the two Gaussian distributions $\mathbb{P} = \mathcal{N}(\mu, \Sigma)$ and $\hat{\mathbb{P}} = \mathcal{N}(\hat{\mu}, \hat{\Sigma})$. The squared 2-Wasserstein distance between them can be computed by:

$$\mathbb{W}^2(\hat{\mathbb{P}}, \mathbb{P}) = \|\mu - \hat{\mu}\|_2^2 + \text{tr} \left(\Sigma + \Sigma_0 - 2(\hat{\Sigma}^{1/2} \Sigma \hat{\Sigma}^{1/2})^{1/2} \right).$$

By $\mathbb{W}(\hat{\mathbb{P}}, \mathbb{P}) \leq r$, we have

$$\begin{aligned} \|\mu - \hat{\mu}\|_2^2 &\leq r, \\ \text{tr} \left(\Sigma + \Sigma_0 - 2(\hat{\Sigma}^{1/2} \Sigma \hat{\Sigma}^{1/2})^{1/2} \right) &\leq r. \end{aligned}$$

By the Procrustes representation of the Bures metric [4], we have

$$\text{tr} \left(\Sigma + \Sigma_0 - 2(\hat{\Sigma}^{1/2} \Sigma \hat{\Sigma}^{1/2})^{1/2} \right) = \min_{U^\top U = I} \|\Sigma^{1/2} - \hat{\Sigma}^{1/2} U\|_F,$$

where $\|\cdot\|_F$ is the Frobenius norm, i.e., $\|A\|_F = \sqrt{\text{tr}(A^\top A)}$. Therefore, we know that, there exists a unitary matrices U^* such that

$$\|\Sigma^{1/2} - \hat{\Sigma}^{1/2} U^*\|_F \leq r.$$

Since the spectral norm is bounded by the Frobenius norm, i.e., $\|A\|_2 \leq \|A\|_F$, we have

$$\|\Sigma^{1/2} - \hat{\Sigma}^{1/2} U^*\|_2 \leq r.$$

From $\|\hat{\Sigma}^{1/2} U^*\|_2 = \|\hat{\Sigma}^{1/2}\|_2 = \sqrt{\lambda_{\max}(\hat{\Sigma})}$, we have

$$\|\Sigma^{1/2}\|_2 \leq \|\hat{\Sigma}^{1/2} U^*\|_2 + \|\Sigma^{1/2} - \hat{\Sigma}^{1/2} U^*\|_2 \leq \sqrt{\lambda_{\max}(\hat{\Sigma})} + r.$$

Squaring both sides yields

$$\lambda_{\max}(\Sigma) \leq (\sqrt{\lambda_{\max}(\hat{\Sigma})} + r)^2 = \bar{\lambda}.$$

For the lower bound, we apply the standard singular value perturbation bound [21], which states that for any matrices A and B , we have

$$\sigma_{\min}(A) \geq \sigma_{\min}(B) - \|A - B\|_2.$$

Setting $A = \Sigma^{1/2}$ and $B = \hat{\Sigma}^{1/2} U^*$:

$$\sigma_{\min}(\Sigma^{1/2}) \geq \sigma_{\min}(\hat{\Sigma}^{1/2} U^*) - r = \sqrt{\lambda_{\min}(\hat{\Sigma})} - r$$

Since the singular value cannot be negative, we have:

$$\sqrt{\lambda_{\min}(\Sigma)} \geq \max \left\{ 0, \sqrt{\lambda_{\min}(\hat{\Sigma})} - r \right\}.$$

Squaring this gives $\lambda_{\min}(\Sigma) \geq \underline{\lambda}$, which completes the proof. \square

With the above over-approximation, we replace the distributionally robust constraint in (11) with its surrogate as

$$\inf_{\mathbb{P} \in \mathbb{D}_{sur}(t)} \mathbb{P}(G(u(t), \xi(t), v_\xi(t)) \leq 0) \geq 1 - \epsilon.$$

Furthermore, to obtain a convex sufficient condition for the worst-case chance constraint, we upper bound the tail risk via the Conditional Value-at-Risk (CVaR) [18] as follows

$$\sup_{\mathbb{P} \in \mathbb{D}_{sur}(t)} \text{CVaR}_{1-\epsilon}^{\mathbb{P}} [G(u(t), \xi(t), v_\xi(t))] \leq 0. \quad (15)$$

Next, we specialize the above worst-case CVaR constraint to the concrete CBF condition in (2) and derive a tractable reformulation that can be enforced online by the worst-case bounding over the surrogate ambiguity set. We summarize our main result of this paper as the following theorem.

Theorem 1. Consider the safety function in the form of (2) and choose the class- \mathcal{K} function as $\alpha(x) = \alpha x$. Let $u(t)$ be an optimal solution to the following tractable convex optimization problem:

$$\min_{u \in \mathbb{R}^m} \|u - u_{norm}\|_2^2 \quad (16a)$$

$$\text{s.t. } \beta \geq \|a(x, u)\|_2^2 \quad (16b)$$

$$a(x, u) = f(x) + g(x)u \quad (16c)$$

$$\begin{aligned} 2a^\top(\hat{\mu}_\xi - x) + (2r + 2\tau(\epsilon)\sqrt{\bar{\lambda}}) \beta + \alpha d_s^2 \\ + \frac{1}{\alpha\epsilon} (\|\hat{\mu}_v\|_2 + r + n\bar{\lambda}) \leq 0 \end{aligned} \quad (16d)$$

where $\hat{\mu}_\xi(t)$ and $\hat{\mu}_v(t)$ denote the estimates of the obstacle position and velocity obtained from the Kalman filtering, respectively. Then the resulting control input $u(t)$ satisfies the distributionally robust safety constraint (15).

Proof. Consider the safety function in (2). The corresponding constraint function can be written as

$$G(u, \xi, v_\xi) = -2(x - \xi)^\top a - 2(x - \xi)^\top v_\xi - \alpha(\|x - \xi\|_2^2 - d_s^2).$$

To handle the coupling term $2(\xi - x)^\top v_\xi$, we apply Young's inequality: for any $\eta > 0$,

$$2(\xi - x)^\top v_\xi \leq \frac{1}{\eta} \|x - \xi\|_2^2 + \eta \|v_\xi\|_2^2.$$

Choosing $\eta = \frac{1}{\alpha}$ yields

$$G(u, \xi, v_\xi) \leq -2(x - \xi)^\top a(x, u) + \frac{1}{\alpha} \|v_\xi\|_2^2 + \alpha d_s^2.$$

Therefore, we have

$$\begin{aligned} \text{CVaR}_{1-\epsilon}^{\mathbb{P}} [G(u, \xi, v_\xi)] &\leq \text{CVaR}_{1-\epsilon}^{\mathbb{P}} [-2(x - \xi)^\top a(x, u)] \\ &\quad + \text{CVaR}_{1-\epsilon}^{\mathbb{P}} \left[\frac{1}{\alpha} \|v_\xi\|_2^2 \right] + \alpha d_s^2 \end{aligned} \quad (17)$$

Since $\xi \sim \mathcal{N}(\mu_\xi, \Sigma_{\xi\xi})$, where $\Sigma_{\xi\xi}$ denotes the position covariance, i.e., the position block of the obstacle-state covariance matrix, it follows that the scalar random variable $2a^\top(\xi - x)$ is a Gaussian, i.e.,

$$2a^\top(\xi - x) \sim \mathcal{N}(2a^\top(\mu_\xi - x), 4a^\top \Sigma_{\xi\xi} a).$$

Therefore, by the closed-form expression of CVaR for a univariate Gaussian random variable [18], we have

$$\text{CVaR}_{1-\epsilon}^{\mathbb{P}} (-2(x - \xi)^\top a) = 2a^\top(\mu_\xi - x) + 2\sqrt{a^\top \Sigma_{\xi\xi} a} \cdot \tau(\epsilon)$$

where

$$\tau(\epsilon) = \frac{\varphi(\Phi^{-1}(1 - \epsilon))}{\epsilon}$$

and φ and Φ are the probability density function and cumulative distribution function, respectively.

Since $\mathbb{P} \in \mathbb{D}_{\text{sur}}$ by Proposition 1, we have $\|\mu_\xi - \hat{\mu}_\xi\|_2 \leq r$ and $\Sigma_{\xi\xi} \preceq \bar{\lambda}I$. Hence it holds that

$$\begin{aligned} 2a^\top(\mu_\xi - x) &= 2a^\top(\hat{\mu}_\xi - x) + 2a^\top(\mu_\xi - \hat{\mu}_\xi) \\ &\leq 2a^\top(\hat{\mu}_\xi - x) + 2r\|a\|_2, \end{aligned} \quad (18)$$

Moreover, we have

$$a^\top \Sigma_{\xi\xi} a \leq a^\top (\bar{\lambda}I) a = \bar{\lambda}\|a\|_2^2 \Rightarrow 2\sqrt{a^\top \Sigma_{\xi\xi} a} \leq 2\sqrt{\bar{\lambda}}\|a\|_2.$$

Substituting these bounds into the Gaussian CVaR expression yields

$$\text{CVaR}_{1-\epsilon}^{\mathbb{P}}(-2(x-\xi)^\top a) \leq 2a^\top(\hat{\mu}_\xi - x) + (2r + 2\tau(\epsilon)\sqrt{\bar{\lambda}})\|a\|_2.$$

On the other hand, since $\|v_\xi\|_2^2 \geq 0$, we use the standard bound $\text{CVaR}_{1-\epsilon}^{\mathbb{P}}(Z) \leq \frac{1}{\epsilon}\mathbb{E}_{\mathbb{P}}[Z]$ for $Z \geq 0$ to obtain

$$\text{CVaR}_{1-\epsilon}^{\mathbb{P}}(\|v_\xi\|_2^2) \leq \frac{1}{\epsilon}\mathbb{E}_{\mathbb{P}}[\|v_\xi\|_2^2] = \frac{1}{\epsilon}(\|\mu_v\|_2^2 + \text{tr}(\Sigma_{vv})).$$

For $\mathbb{P} \in \mathbb{D}_{\text{sur}}$, we have $\|\mu_v - \hat{\mu}_v\|_2 \leq r$ and $\Sigma_{vv} \preceq \bar{\lambda}I$, which imply $\|\mu_v\|_2 \leq \|\hat{\mu}_v\|_2 + r$ and $\text{tr}(\Sigma_{vv}) \leq n\bar{\lambda}$, where n is the dimension of v_ξ . Therefore,

$$\text{CVaR}_{1-\epsilon}^{\mathbb{P}}\left(\frac{1}{\alpha}\|v_\xi\|_2^2\right) \leq \frac{1}{\alpha\epsilon}((\|\hat{\mu}_v\|_2 + r)^2 + n\bar{\lambda}).$$

Combining the above bounds, we obtain

$$\begin{aligned} \text{CVaR}_{1-\epsilon}^{\mathbb{P}}[G(u, \xi, v_\xi)] &\leq 2a^\top(\hat{\mu}_\xi - x) + (2r + 2\tau(\epsilon)\sqrt{\bar{\lambda}})\|a\|_2 \\ &\quad + \alpha d_s^2 + \frac{1}{\alpha\epsilon}((\|\hat{\mu}_v\|_2 + r)^2 + n\bar{\lambda}). \end{aligned}$$

By introducing an auxiliary variable β to lower bound $\|a(x, u)\|_2$, i.e., $\|a(x, u)\|_2 \geq \beta$, the above constraint gives a tractable convex constraint, which completes the proof. \square

Remark 1. *The optimization problem in (16) is a second-order cone program (SOCP) and is therefore tractable. The objective $\|u - u_{\text{norm}}\|_2^2$ is a convex quadratic function. Moreover, $a(x, u) = f(x) + g(x)u$ is affine in u , and the constraint linking β and $\|a(x, u)\|_2$ can be expressed as a second-order cone constraint. All remaining terms in the robust safety constraint are affine in the decision variables (u, β) , since $\hat{\mu}_\xi(t)$, $\hat{\mu}_v(t)$, $r(t)$, $\bar{\lambda}(t)$, α , and ϵ are known at time t . Consequently, the feasible set is convex and the problem can be solved efficiently using standard SOCP solvers, enabling real-time controller synthesis.*

V. EXPERIMENTAL VALIDATION

In this section, we validate the effectiveness and robustness of the proposed Distributionally Robust Control Barrier Function (DR-CBF) framework through simulation studies. We focus on safe task execution for robotic manipulators operating in environments with dynamic obstacles and perception uncertainty. The performance of our approach is evaluated against baseline methods.

TABLE I: Statistical Results of 1000 Trials

Environment Noise	Method	Success Rate	Collision Rate	Stuck Rate
$\sigma = 0.01$	Nominal-CBF	0.868	0.131	0
	DR-CBF	0.968	0.032	0
$\sigma = 0.05$	Nominal-CBF	0.699	0.301	0
	DR-CBF	0.979	0.210	0
$\sigma = 0.1$	Nominal-CBF	0.483	0.517	0
	DR-CBF	0.828	0.090	0.082

A. Experimental Setup

We consider a robotic manipulator operating in a workspace with dynamic obstacles, where the robot arm must navigate from a starting position to a target goal while avoiding a **moving obstacle**. The state of the obstacle is subject to varying levels of environmental noise, modeled as Gaussian noise with standard deviation $\sigma \in \{0.01, 0.05, 0.1\}$.

We conducted 1000 simulation trials for each noise level and measured the following metrics:

- **Success Rate:** The percentage of trials where the agent successfully reaches the goal without any collision;
- **Collision Rate:** The percentage of trials resulting in a collision with the obstacle;
- **Stuck Rate:** The percentage of trials where the agent fails to reach the goal within a predefined time limit due to overly conservative behavior.

B. Statistical Results

The statistical results of the comparative study are summarized in Table I. As shown in Table I, when the noise level is low ($\sigma = 0.01$), both methods achieve relatively high success rates. However, as the noise intensity increases to $\sigma = 0.1$, the success rate of the nominal CBF drops significantly to 48.3%, with a collision rate exceeding 50%.

In contrast, our DR-CBF maintains a high success rate of 82.8% even under extreme noise ($\sigma = 0.1$). Although it exhibits a small stuck rate (8.2%), it reduces the collision rate from 51.7% to 9%. These results show that our distributionally robust approach effectively mitigates the ‘‘reality gap’’ between the estimated and true obstacle distributions.

C. Qualitative Results and Visualization

To provide further intuition into the safety-critical behavior of our algorithm, we visualize the trajectories and the adaptive safety boundaries in Fig 1.

As illustrated in Fig 1, where the solid red circle represents the true position of the dynamic obstacle, and the dashed yellow circle denotes the obstacle position estimated by the Kalman filter. In nominal CBF case (Fig 1(a)), the agent enters a collision course because the estimated position deviates from the truth. Conversely, our DR-CBF (Fig 1(b)) successfully avoids the hazard by employing an uncertainty-aware safety constraint. Note that when the prediction error is small, both methods successfully avoid obstacles; however, when the prediction error is large, nominal CBF fails to avoid obstacles while DR-CBF succeeds, as shown in Fig 3.

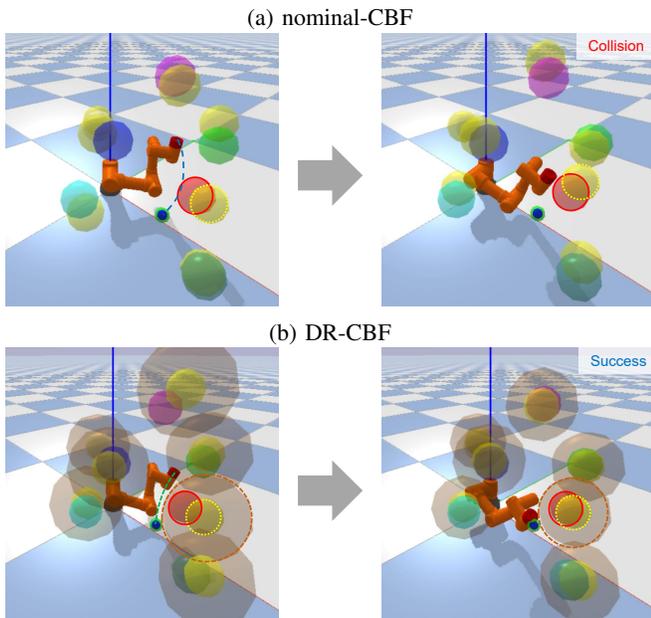


Fig. 1: Example trials of the nominal-CBF and our proposed DR-CBF methods in an uncertain perception scenario.

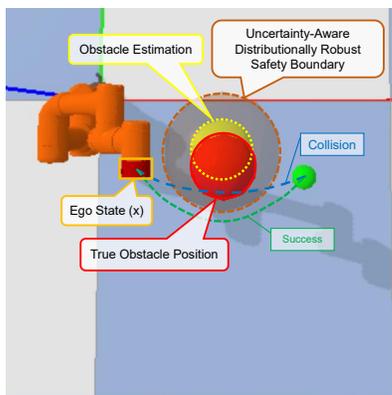


Fig. 2: Illustration of our proposed DR-CBF method

The core mechanism of our algorithm is visualized in Fig 2. Based on the proposed distributionally robust framework, we construct a “surrogate ambiguity set” that encapsulates the true obstacle distribution within a Wasserstein ball. This manifests as a larger safety ball (the DRO safety boundary) around the estimated position. Importantly, this safety boundary is adaptive: as the perception noise σ increases (as seen in the comparison between Fig 4(a), (b) and (c)), the Kalman filter’s covariance grows, causing the DRO algorithm to automatically expand the safety radius to ensure a higher probability of safety. This adaptive conservativeness allows the agent to maintain high efficiency when the perception is confident, while enforcing strict safety margins when uncertainty is high.

In addition, we present plots illustrating the distances to the surfaces of individual obstacles during the robotic manipulators obstacle avoidance process, as shown in Fig 5. Note that, given the robotic manipulators’ multi-link structure,

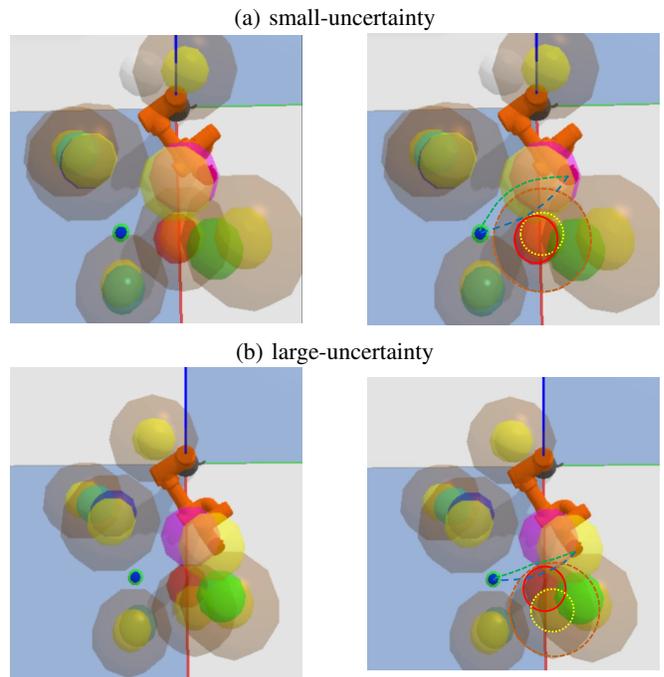


Fig. 3: Comparison in (a) small uncertainty and (b) large uncertainty perception scenarios.

each plotted curve tracks the minimum distance to the obstacle among all links over time.

VI. CONCLUSION

This paper studied safe control of robotic manipulators operating in dynamic environments under perception uncertainty. We proposed a distributionally robust CBF framework that integrates Kalman-filter-based obstacle estimation with a Wasserstein ambiguity set to explicitly account for the mismatch between the estimated and true obstacle-state distributions. To enable real-time implementation, we introduced a surrogate ambiguity set that yields a conservative but tractable convex reformulation, leading to an online SOCP-based safety filter with worst-case CVaR guarantees. Simulation results demonstrated that the proposed DR-CBF significantly reduces collision rates and improves success rates compared to nominal CBF-based controller, especially under high measurement noise. Future work will focus on extending the framework to more general belief representations as well as validating the approach on real robotic hardware with vision-based perception in cluttered dynamic scenes.

REFERENCES

- [1] Aaron Ames, Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. Control barrier functions: Theory and applications. In *European Control Conference (ECC)*, pages 3420–3431, 2019.
- [2] Aaron D Ames, Xiangru Xu, Jessy W Grizzle, and Paulo Tabuada. Control barrier function based quadratic programs for safety critical systems. *IEEE Transactions on Automatic Control*, 62(8):3861–3876, 2017.
- [3] Liviu Aolaritei, Marta Fochesato, John Lygeros, and Florian Dörfler. Wasserstein tube MPC with exact uncertainty propagation. In *IEEE Conference on Decision and Control (CDC)*, pages 2036–2041, 2023.

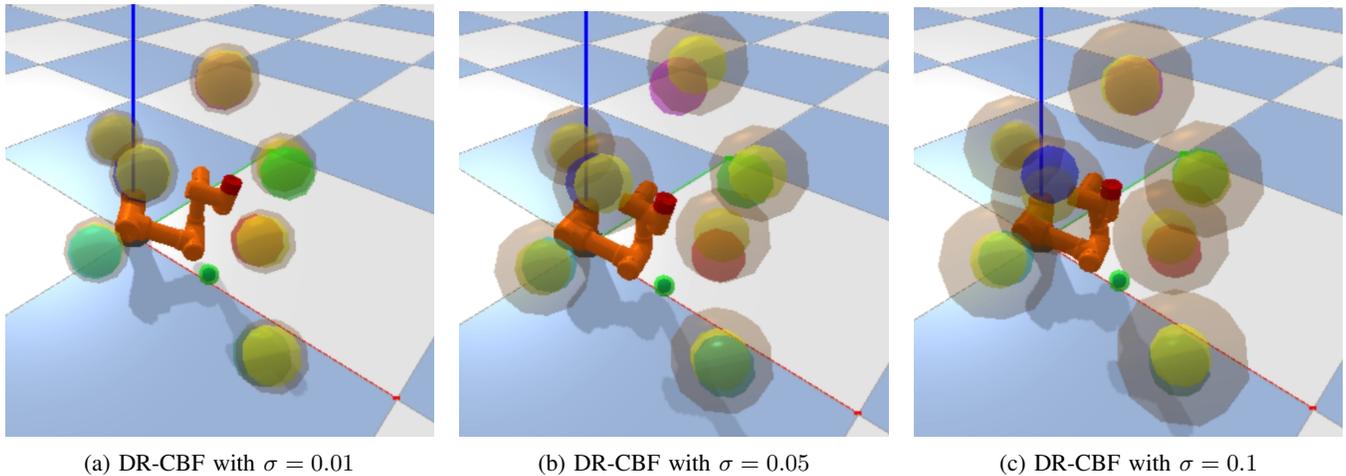
(a) DR-CBF with $\sigma = 0.01$ (b) DR-CBF with $\sigma = 0.05$ (c) DR-CBF with $\sigma = 0.1$

Fig. 4: The size of the safety boundary under different noise intensities

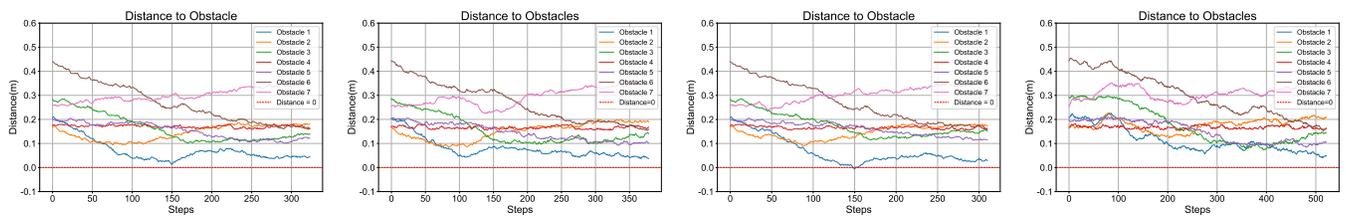
(a) nominal-CBF with $\sigma = 0.05$.(b) DR-CBF with $\sigma = 0.05$.(c) nominal-CBF with $\sigma = 0.1$.(d) DR-CBF with $\sigma = 0.1$.

Fig. 5: Distance to Obstacle for different CBF methods and noise levels.

- [4] Rajendra Bhatia, Tanvi Jain, and Yongdo Lim. On the bures–wasserstein distance between positive definite matrices. *Expositiones mathematicae*, 37(2):165–191, 2019.
- [5] Jeremy Coulson, John Lygeros, and Florian Dörfler. Distributionally robust chance constrained data-enabled predictive control. *IEEE Transactions on Automatic Control*, 67(7):3289–3304, 2021.
- [6] Ersin Das and Joel W Burdick. Robust control barrier functions using uncertainty estimation with application to mobile robots. *arXiv preprint arXiv:2401.01881*, 2024.
- [7] Vikas Dhiman, Mohammad Javad Khojasteh, Massimo Franceschetti, and Nikolay Atanasov. Control barriers in bayesian learning of system dynamics. *IEEE Transactions on Automatic Control*, 68(1):214–229, 2023.
- [8] Yousef Emam, Paul Glotfelter, Sean Wilson, Gennaro Notomista, and Magnus Egerstedt. Data-driven robust barrier functions for safe, long-term operation. *IEEE Transactions on Robotics*, 38(3):1671–1685, 2022.
- [9] Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations. *Mathematical Programming*, 171:115–166, 2018.
- [10] D. Fox, W. Burgard, and S. Thrun. The dynamic window approach to collision avoidance. *IEEE Robotics Autom. Mag.*, 4:23–33, 1997.
- [11] Mikhail Koptev, Nadia Figueroa, and Aude Billard. Neural joint space implicit signed distance functions for reactive robot manipulator control. *IEEE Robotics and Automation Letters*, 8(2):480–487, 2023.
- [12] Paul Lathrop, Beth Boardman, and Sonia Martínez. Distributionally safe path planning: Wasserstein safe RRT. *IEEE Robotics and Automation Letters*, 7(1):430–437, 2021.
- [13] Yiming Li, Yan Zhang, Amirreza Razmjoo, and Sylvain Calinon. Representing robot geometry as distance fields: Applications to whole-body manipulation. In *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, pages 15351–15357, 2024.
- [14] Kehan Long, Yin Zhuang Yi, Jorge Cortés, and Nikolay Atanasov. Safe and stable control synthesis for uncertain system models via distributionally robust optimization. In *American Control Conference (ACC)*, pages 4651–4658, 2023.
- [15] Kehan Long, Yin Zhuang Yi, Zhirui Dai, Sylvia Herbert, Jorge Cortés, and Nikolay Atanasov. Sensor-based distributionally robust control for safe robot navigation in dynamic environments. *The International Journal of Robotics Research*, 45(2):328–351, 2026.
- [16] Helen Oleynikova, Zachary Taylor, Marius Fehr, Roland Siegwart, and Juan Nieto. Voxblox: Incremental 3d Euclidean signed distance fields for on-board MAV planning. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1366–1373, 2017.
- [17] Allen Z. Ren and Anirudha Majumdar. Distributionally robust policy learning via adversarial environment generation. *IEEE Robotics and Automation Letters*, 7(2):1379–1386, 2022.
- [18] R. Tyrrell Rockafellar and Stanislav Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26(7):1443–1471, 2002.
- [19] Kanghyun Ryu and Negar Mehr. Integrating predictive motion uncertainties with distributionally robust risk-aware control for safe robot navigation in crowds. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2410–2417, 2024.
- [20] Simon Schaefer, Karen Leung, Boris Ivanovic, and Marco Pavone. Leveraging neural network gradients within trajectory optimization for proactive human-robot interactions. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9673–9679. IEEE, 2021.
- [21] Gilbert W Stewart and Ji-guang Sun. Matrix perturbation theory. (*No Title*), 1990.
- [22] Bart P. G. Van Parys, Daniel Kuhn, Paul J. Goulart, and Manfred Morari. Distributionally robust control of constrained stochastic systems. *IEEE Transactions on Automatic Control*, 61(2):430–442, 2016.
- [23] Yujie Wang and Xiangru Xu. Disturbance observer-based robust control barrier functions. In *American Control Conference (ACC)*, pages 3681–3687, 2023.
- [24] Weijun Xie. On distributionally robust chance constrained programs with Wasserstein distance. *Math. Program.*, 186:115–155, 2021.